The Grid: An IT Infrastructure for NOAA in the 21st Century

Mark W. Govett[1], Mike Doney and Paul Hyder[2]
NOAA Forecast Systems Laboratory
325 Broadway, Boulder Colorado 80305
May 2004

EXECUTIVE SUMMARY

This paper highlights the need to build a grid infrastructure to meet the challenges facing NOAA in the 21st century. Given the enormous expected demands for data, and increased size and density of observational systems, current systems will not be scalable for future needs without incurring enormous costs. NOAA's IT infrastructure is currently a set of independent systems that have been built up over time to support its programs and requirements. NOAA needs integrated systems capable of handling a huge increase in data volumes from expected launches of GOES-R, NPOESS, new observing systems being proposed or developed, and to meet requirements of the Integrated Earth Observation System. Further, NOAA must continue moving toward integrated compute resources to reduce costs, to improve systems utilization, to support new scientific challenges and to run and verify increasingly complex models using next generation high-density data streams. Finally, NOAA needs a fast, well-managed network capable of meeting the needs of the organization: to efficiently distribute data to users, to provide secure access to IT resources, and be sufficiently adaptable and scalable to meet unanticipated needs in the future.

Just as the NOAA Observing System Architecture (NOSA) was developed to define, coordinate and integrate data sources, NOAA needs to build an IT counterpart (NITSA) to define, coordinate and integrate its IT resources. We propose the construction of a grid infrastructure, centered around a fast network, which will permit resources including computers, data storage, software systems and services to be managed and shared more effectively. All three types of grids identified in this paper, can be utilized at NOAA: compute grids to provide access to under-utilized cycles and to link super-computing centers, data grids to promote sharing and better utilization of data, and service grids to provide the reliability and redundancy demanded by operational processes. Grid extends the power and success of the World Wide Web (WWW) by providing robust mechanisms for data access, discovery and integration that can transform static methods of data distribution into dynamic demand driven delivery systems. We have also described a process by which secure grids can be built using token-based authentication, VPN communications between secure NOAA systems, and restricted external access through heavily monitored firewalls.

NOAA currently spends hundreds of millions of dollars on observing systems, supercomputing facilities, data centers, and the network. The development of a NITSA, enabled by Grid, can potentially reap huge cost savings for the organization. We describe the value of building an efficient, scalable, secure and integrated network. It is not sufficient or cost effective to continue

---

[1] *Email Address: Mark.W.Govett@noaa.gov*
[2] *Cooperative Institute for Research in the Atmosphere, Colorado State University, Fort Collins, CO 80523 USA*

to upgrade select point-to-point network links. The entire network needs to be examined and an integrated, managed, coordinated network resource must be created that is an asset to the entire organization. We also describe the value of building an grid infrastructure to coordinate access to the IT resources necessary to handle increasing volumes of data, to develop and run next generation models and to evaluate the cost /benefit of new and existing observing systems. As NOAA stresses the increased importance of sharing IT resources in the interest of cost savings, code interoperability across HPC systems at GFDL, FSL and NCEP has become a critical requirement. It is not enough to simply achieve code interoperability, however; improved usability of these codes is also required. Grids provide a supportive distributed computing environment in which to run portable codes and simplifies user access to shared system resources. Finally, we describe the importance of VOs (Virtual Organizations) to improve collaborations between organizational units, to improve efficiency and reduce costs. Grids support the creation of VOs, which can map IT resources directly into cross cutting programs identified by NOAA's PPBES management plan.

The Grid: An IT Infrastructure for NOAA in the 21st Century

Mark W. Govett, Mike Doney and Paul Hyder
NOAA Forecast Systems Laboratory
325 Broadway, Boulder Colorado 80305
May 2004

1. Introduction

The NOAA Strategic Plan, titled "New Priorities for the 21st Century," highlights many significant and complex challenges facing NOAA. The strategic plan states: "NOAA's role is to predict environmental changes, protect life and property, provide decision makers with reliable scientific information, manage the Nation's living marine and coastal resources, and foster global environmental stewardship" [2003]. NOAA's mission is broad, ranging from managing coastal and marine resources to predicting changes in the Earth's environment. Data are central to this mission including, (1) building and maintaining observational platforms such as radars, satellite-based instruments, aircraft, profilers, weather balloons, ships, and buoys, (2) creating data products for environmental prediction, (3) disseminating the data to users, and (4) archiving the data so they can be used by research laboratories to develop new capabilities and forecast products.

NOAA has made a significant investment in Information Technology (IT) for computers, data storage and networks required to handle the processing, dissemination, archival, and management of its data. NOAA IT assets include three super-computing centers maintained at the Forecast Systems Laboratory (FSL), Geophysical Fluid Dynamics Laboratory (GFDL), and at the National Centers for Environmental Prediction (NCEP). Compute resource also include thousands of support systems ranging from desktop systems to high performance servers used for data processing, new product development and ongoing scientific research. NCDC, NGDC and NODC are NOAA's primary facilities for the long-term archival of data, but many other secondary facilities store data of interest to their local facilities for shorter periods. Mass storage requirements for short and long term data storage are significant. To effectively receive and disseminate these data, NOAA also relies heavily on a network composed of NOAAPORT, the Internet, and many dedicated point-to-point connections.

NOAA's IT investment directly relates to the volume of data it receives, processes and disseminates. Data volume has grown tremendously over the last decade. Data managers estimate NOAA will archive more data in 2004 than is contained in its entire archive through 1998. This volume of data handling requires large investments in facilities, hardware, software and staff. For example, NESDIS recently began construction on a $61 million facility designed to handle ingesting and processing of high-volume next-generation satellite data. NCEP recently signed a $240 million 9 year contract with IBM to provide supercomputing facilities for NWS operations. This contract includes computing, mass storage, and the housing and maintenance of a system that is composed of two virtual systems: an operational platform and a research and development platform. There will also be a backup system to the operational platform in FY05.

The huge growth in data volume is expected to continue. In the next 10 years, data growth is expected to increase by more than 100 times over current volumes. The number of weather satellites is expected to increase from a few dozen to more than 200 containing 300 instruments in the next 15 years. New satellite technology offers the most compelling glimpse into NOAA's ingest, dissemination and archival needs for the future. The National Polar-orbiting Operational Environmental Satellite System (NPOESS) program is a 6 billion dollar joint NOAA and DoD program that will begin launching a series of polar orbiters beginning in 2009. The GOES-R satellites are slated to be online in 2012 and will record huge increases in data over existing GOES platforms. Effective use of this satellite data will depend on sufficient IT infrastructure to ingest, process, distribute and archive the data. This IT infrastructure is costly both in terms of increasing the capacity of systems, storage, and networks, and requiring sufficient staff to support the development and maintenance of their hardware and software systems. Better coordination and integration of NOAA's IT assets could lead to significant cost reductions and better utilization of systems, software and staff.

The need to better coordinate NOAA's IT assets will be required to meet the expected demands of the proposed Integrated Earth Observation System. Since 2003, participants have been meeting to promote and design a comprehensive, coordinated, and sustained Earth observation system or systems among governments and the international community to understand and address global environmental and economic challenges (www.earthobservationsummit.gov). While the importance of developing such a system is clear, the task of building it is daunting. The size of the proposed system indicates it will need to be distributed across nations and mechanisms will need to be developed to share these data effectively. NOAA's role in developing this system will be pivotal to its success.

To respond to these changing requirements and mission goals, NOAA's IT infrastructure must become increasingly robust, responsive, flexible, efficient, secure, adaptable and cost-effective. Grid computing has emerged as a viable technology that is increasingly utilized by government laboratories, corporate institutions, and high performance computing centers around the world. Grid computing is being used in diverse applications such as high-energy physics, medical imaging, meteorology, and business applications. Advancements in high speed networks, the development of software to enable efficient, secure distributed computing, and the role of industry in supporting and promoting Grid, have moved grid computing from a fringe technology just a decade ago, to an increasingly mainstream technology today.

Grid computing is being explored at the Forecast Systems Laboratory as a means to provide better integration of NOAA's growing data, compute and network resources and to meet the expected needs of these and other projects in the future. Specifically, grid computing technologies can be used to (1) provide better utilization of NOAA's IT assets, (2) provide a more efficient way to ingest, disseminate and archive data to NOAA constituents and the general public, (3) reduce costly duplication in data archival and product generation systems and the network, (4) improve the utilization of data handled by NOAA, and (5) provide more efficient access to compute facilities and requisite storage facilities. A general description of grid computing is given in Section 2. Section 3 shows how grids can benefit NOAA and highlights the need to build a more robust, cost-effective NOAA network. Section 4 identifies several challenges to developing grids, highlights FSL's progress in exploring Grid, and proposes a process to begin building and deploying grids at NOAA.

2. Grid Computing

The modern concept of grid computing began in the mid-1990s when the term "the Grid" was coined to describe a computing infrastructure to support large-scale resource sharing necessary for scientific research. Software development efforts began in earnest, and in 2000 the first of two foundational papers was written by Foster, et al., to describe functional requirements for a grid computing infrastructure and the development of the Globus software intended to meet those requirements. The first paper, titled "The Anatomy of the Grid", defined essential properties of Grids to build Virtual Organizations (VOs) and proposed a set of protocols and services to provide access, coordination and sharing of distributed resources [Foster et al. 2001]. A VO is defined to be those distributed communities willing to share their resources in order to achieve common goals. The second paper, titled "The Physiology of the Grid", describes the interactions between grid components and defines a set of protocols required for interoperability between them [Foster, et al., 2002]. The central tenet of these papers is to move toward an end goal of ubiquitous computing serving the needs of VOs, where computing is available on-demand and users do not have to be concerned with where their tasks are running or where data reside.

Industry support for grid computing has coalesced around Globus, a software toolkit designed to handle resource sharing and discovery, data movement, security, access and other operations required in distributed computing environments. Recently completed efforts to define grid software standards that align with standards from the World Wide Web, and Globus's adherence to these standards, have strengthened the interest and commitment of industry toward grid computing. With support from industry heavyweights such as IBM, HP, Microsoft, and Sun, the Globus toolkit and related grid standards represent a roadmap to generalized grid computing in the future.

The most common analogy for grid computing is the electric power grid, a system of power generating plants connected together via transmission lines that supply power to the grid for use by consumers on demand. Instead of a system of individual community power plants, the development of a power grid has increased the availability and reliability of power. As a measure of its success, consumers no longer think in terms of how power gets to their homes, it is simply available to them whenever they need it. Key to the success of the power grid has been the development and adherence to power standards (eg. voltage, amps, cycles), which permit consumers to access and use the resources in the same way anywhere on the grid [Wladawsky-Berger 2004].

Similarly, grid computing, or simply "the grid" is most often described in terms of computer hardware: typically compute resources, data storage, and a high-speed network to link them together into a single IT entity. Grid resources can be located at a single site or distributed across the country or around the world. They can be confined to a single institution or include multiple collaborating organizations. They can be restricted to a simple network of desktop systems, or include multiple supercomputing sites, mass storage facilities, and dedicated high speed wide-area networks.

From the user perspective, the grid looks like a single large IT resource with a single sign-on or login required. Users obtain a grid account from their organization and log on to the grid using standardized authentication and access procedures. Globus uses the X.509 standard for remote access, the same security model with public and private encryption used by ssh (secure shell), the preferred method of secure remote access used on most systems today. Once on the grid, users have access to all grid resources they are authorized to use; they do not need to log-on to each system where access is required. This reduces the complexity of managing individual accounts on multiple systems where users must remember different system access procedures, system administrators must maintain individual accounts on each system, and security personnel must monitor remote access between systems.

The Globus toolkit is middleware designed to mediate the complexity of providing access to distributed hardware resources, from the user applications that need them. Figure 1 illustrates the four layers of grid: the hardware infrastructure, the Globus toolkit, grid tools and grid-enabled user applications. The Globus and grid tools layers will be described in more detail.

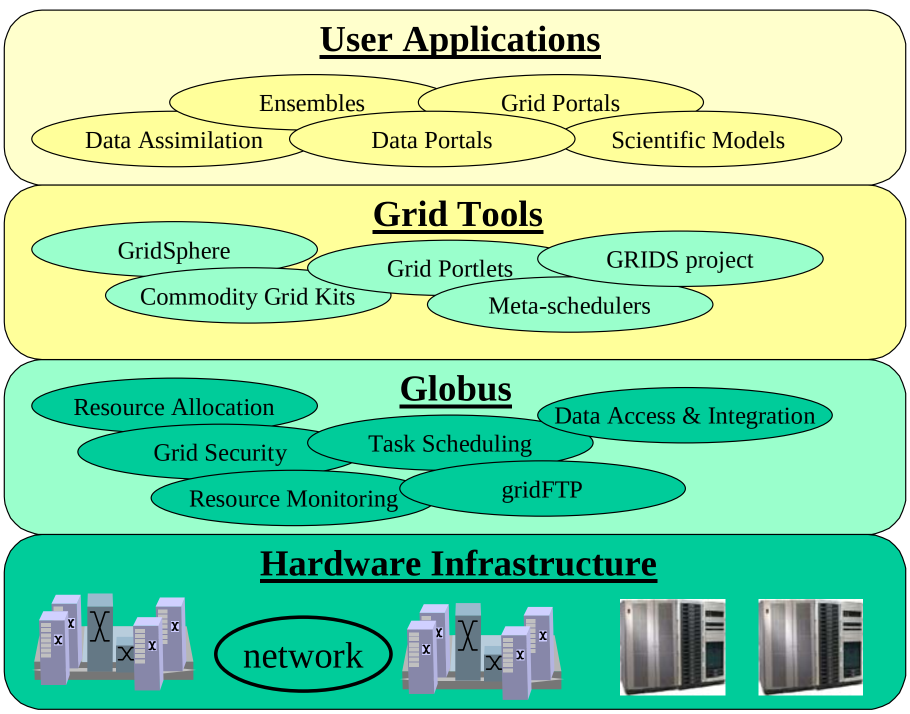

Figure 1: Globus was designed to serve as middleware, mediating the complexity of providing access to distributed and shared hardware resources to applications that require them.

2.1 The Globus Toolkit

The Globus development represents a successful partnership between research and industry and includes contributions from the Globus Alliance, Industry and the Globus Grid Forum (GGF). The Globus Alliance, responsible for the development and free support of Globus has four core participants: Argonne National Laboratory, University of Southern California, University of Edinburgh, and UK E-science. Industry support for Globus include partners such as IBM, HP, Sun and Platform Computing who have contributed to the development both of grid standards and the software. Several companies also provide commercial support for Globus, including IBM, Platform Computing, and Data Synapse; other vendors have plans to add support for Globus in the near future. Oracle, for example, has assigned a large staff to build “grid enabled” database applications, and to support the expected needs of the business and scientific communities.

The GGF, represented by both industry and research groups, is primarily concerned with the development of Grid standards. These standards provide common application programming interfaces (APIs) upon which Globus is based, and permit others, including industry, to contribute to and support the toolkit. Globus is standards based software that builds on three web service standards defined within the W3C (www.w3c.org) and other standards bodies: (1), the Web Services Definition Language (WSDL) provides a means to define a service, (2) Web Service Inspection (WS-Inspection) provides a set of conventions to discover data services, and (3), the Simple Object Access Protocol (SOAP) provides a means for service provider and service requestor to communicate. Grid extends these web services to include grid services and standards for security, resource management, resource discovery, and data access that are briefly described.

2.1.1 Security

System security is critical everywhere but is particularly important in large distributed and diverse grid environments. The grid standard that governs security is called the Grid Security Infrastructure (GSI). Four important requirements in grid security are single sign-on, authentication, delegation, and authorization. Single sign-on is the ability to authenticate on a single source host and to have that sign-on be inherited by jobs running across the grid. Authentication verifies the user’s identity. Typically a user ID and password assigned on the multitude of local hosts is validated against the certificate as part of the grid job initialization. Delegation is the ability to permit another program to act on the user's behalf - in effect giving that program the ability to access all resources that the user is permitted to access. Authorization is the restriction of the subject to grid resources they are permitted to use and are controlled by the individual sites thus eliminating the need for a centrally managed security system.

GSI addresses these requirements and provides secure communication between elements, and security across organizational boundaries. The implementation of GSI is based on public key encryption, X.509 certificates, and the Secure Sockets Layer (SSL) protocol. A trusted third party Certificate Authority (CA) is used to certify that the link between the public key and user’s identity is valid. Every user and service is identified and authenticated by a certificate that contains:

- subject name, the person or object the certificate represents
- public key belonging to the subject
- identity of the issuing CA
- digital signature of the CA

These certificates provide for mutual authentication as long as the issuing CA or CAs are trusted by all subjects.

2.1.2 Resource Management

Globus includes a set of service components called the Globus Resource Allocation Manager (GRAM) that provides a single standard interface for managing resources on remote systems. Client commands are available for submitting and monitoring remote jobs. In addition, an API is provided that can be used to construct meta-schedulers. Meta-schedulers enable the scheduling of jobs across grid resources. They provide support for advanced reservations, job prioritization, and co-allocation of resources at multiple sites. Examples include ClusterResources SILVER, and meta-schedulers built using the Platform Computing developed Community Scheduler Framework.

2.1.3 Resource Discovery

A grid service extends the concept of a web service to provide information about the state of a service. For example, users of applications will want to know if a particular resource is available, its system characteristics (CPU type, system load, job queue length, job queue type, etc.), usage policies (e.g. only night-time access permitted), resource availability (e.g. data storage, network capacity), and process status (eg. job state, cpu time used). The Grid Resource Information Service (GRIS) runs on grid resources and handles the registration of resource characteristics. The Grid Information Index Server (GIIS) is a user accessible directory server that supports the searching for a resource or by specific resource characteristics. Hierarchies of directory services can also be built to permit local level services (e.g. administrative, security management, load balancing) or to support expanded searches across multiple services.

2.1.4 Data Access

Data access and file transfer are accomplished using a service called GridFTP. GridFTP extends the FTP standard to incorporate security protocols, support for partial file access and high-speed parallel file transfers. Performance and reliability are critical issues when accessing large amounts of data so GridFTP supports parallel data transfers that use multiple TCP streams between the source and destination, and striped data transfers to increase parallelism and to allow data objects to be interleaved across multiple hosts. Increasing the size of TCP windows can also improve the efficiency of large file transfers. GridFTP also includes mechanisms to improve the reliability of data transfers, provides mechanisms to restart data transfers, and offers fault recovery procedures to handle network and server failures. In addition, emerging standards for Data Access and Integration (DAI) will provide techniques to integrate data from diverse sources including DBMSs, XML formatted files and simple binary or flat files.

2.2 Grid Tools

Equally important as the Globus middleware is the set of application-based tools that build on core grid capabilities. The Grid Tools Layer, illustrated in Figure 1, provides high-level tools, programming environments and web service capabilities to simplify the development of grid applications. Two primary languages in grid tools development are:

- The eXtensible Markup Language (XML) provides a flexible, self-describing web-based data format that is queriable and used widely in dynamic web applications.
- The Java language is particularly well suited for interacting with web-based applications and is heavily used.

Using these and other web-based languages, a large, growing and active community of developers is building many open source tools useful in developing grid applications. Two notable developments are the Commodity Grid (CoG) kits (www-unix.globus.org/cog), and GridSphere (www.gridsphere.org). CoG kits provide a language-based framework for grid projects; currently CoGs are available for Java, Python, CORBA, Perl, and Matlab. GridSphere uses both java and XML to construct grid applications, called portlets, to simplify and speedup the development of grid-based web applications.

2.3 Types of Grids

There are three types of grids: compute grids, data grids and service grids.

2.3.1 Compute Grids

Compute grids, the most familiar type of grid, are generally deployed by an organization and often within a single site. Grid standards development and adherence by Globus middleware have permitted more secure, general-purpose grids to be built that span sites and organizations. These generalized types of grids enable the running of applications that require access to large amounts of compute and data resources that would not be possible within a single site. Additionally, compute grids can be constructed to provide better utilization of existing systems. For example, most user desktop systems are utilized during work periods but rarely used at night or on weekends.

Grids can be constructed to utilize these cycles as a shared resource within a work group or VO. Compute grids can also be constructed between large compute facilities to distribute the workload more evenly between sites. System loads at supercomputing sites often relate to the technology refresh cycle. At the beginning of a HPCS procurement, computers are often under-utilized until researchers are able to fully exploit them. At the end of the procurement cycle, HPCSs are often over-utilized. Compute grids can help smooth out these differences by allowing users to obtain cycles from other sites prior to a new procurement and export cycles to other sites after a technology refresh.

2.3.2 Data Grids

Data grids are used to seamlessly access and share data resources across remote systems. Typically, a thin abstraction layer is created between the data resident at each site and what the user or application sees. This allows users to more easily access and manage the data they require, and allows the organization to better manage its data sources. From the user or application perspective, data that may be distributed across multiple file systems and sites is viewed as a single large file system. Data grids can be constructed to permit read-only access to data or they can be viewed as large read/write distributed file systems.

Data grids are simpler to implement than compute grids since access can be more easily controlled and there are fewer security and resource allocation issues. As a result, many successful data grids are in use today. For example, the Earth Science Grid (ESG) is a Department of Energy (DOE) project that permits researchers to share the results of large climate simulations among collaborators within the U.S. and international climate communities. This kind of structure allows users to share research results more fully, and to coordinate further research among collaborators. The ESG and other examples of data grids will be discussed in Section 2.4.

2.3.3 Service Grids

Service grids build on the success of the World Wide Web and the emergence of e-business in a new area of IT called web services. The use of E-business and E-commerce applications have grown exponentially in the last decade because the enterprise community has seen web services as a means to focus on selecting delivery systems on the basis of price/performance and Quality of Service (QoS) requirements rather than being required to invest heavily in a single compute or software platform. Web services and the development of portable platform independent applications permit software re-use, componentization, and integration. Grid services extend the web services concept by providing task and resource management functions across heterogeneous computing environments.

As illustrated in Figure 2, service grids abstract the infrastructure capabilities of data and compute grids into a task-oriented infrastructure that is available to users and other applications. Services insulate applications from the underlying hardware required to accomplish the task and they support portable deployment on multiple platforms. They can be static, created once and always available or dynamic, created to perform some specific task and then destroyed once that task is completed. Services can be simple (authorization, searching, naming, registry) or complex, combining multiple services into a complex service that encompass the comprehensive requirements of an application. For example, a user wishing to run a weather model could utilize a complex service composed of services to: (1) build initial and boundary condition files on one machine, (2) send the results to a service on another system where the model is run, (3) generate forecast products on a third system and (4) visualize the results on a fourth system. Creating services to accomplish common tasks allows an organization to reduce the effort required to develop, port, and maintain duplicate software systems.

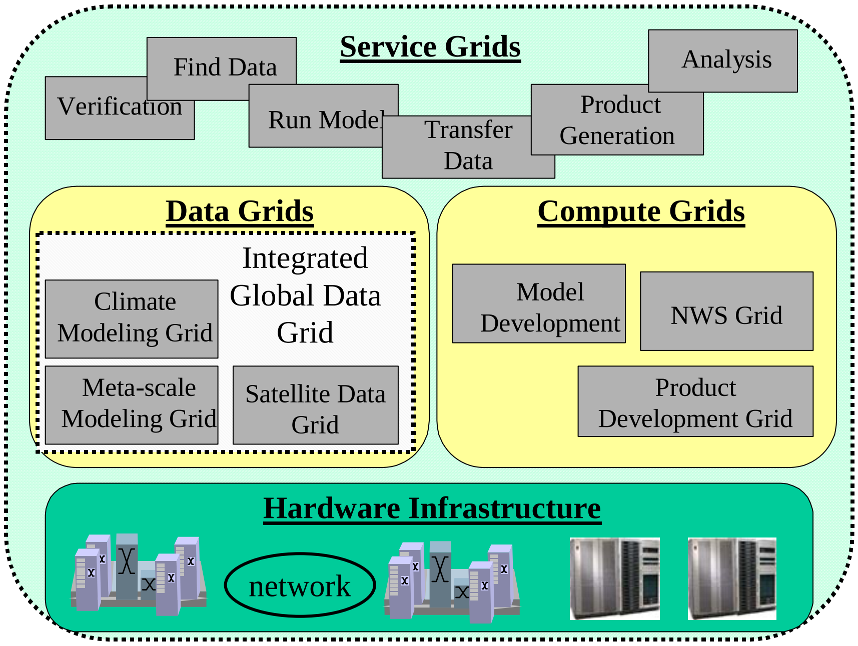

Figure 2: Service grids build on the grid infrastructure provided by compute and data grids in order to accomplish tasks required by users and applications.

2.4 Active Grids

There are hundreds of projects deploying grids to enable new research, and to promote E-business solutions for commercial applications (www-fp.mcs.anl.org/~foster/grid_projects). Three notable grid projects for scientific applications are the TeraGrid, the Earth Science Grid (ESG), and the Linked Environments for Atmospheric Discovery (LEAD).

2.4.1 The TeraGrid

The TeraGrid project, illustrated in Figure 3, was launched by the National Science Foundation in August 2001 with $53 million in funding to four sites: the National Center for Supercomputing Applications (NCSA) at the University of Illinois, Urbana-Champaign, the San Diego Supercomputer Center (SDSC) at the University of California, San Diego, Argonne National Laboratory in Argonne, IL, and Center for Advanced Computing Research (CACR) at the California Institute of Technology in Pasadena. TeraGrid is nearing the end of a multi-year effort to build and deploy the world's largest, most comprehensive, distributed infrastructure for open scientific research. These facilities are capable of managing and storing nearly 1 petabyte of data, high-resolution visualization environments, and provide toolkits for grid computing. Four new TeraGrid sites are expected to add more scientific instruments, large datasets, and additional computing power and storage capacity to the system. All the components are tightly integrated and connected through a fiber-optic network that operates at 40 gigabits per second (www.teragrid.org).

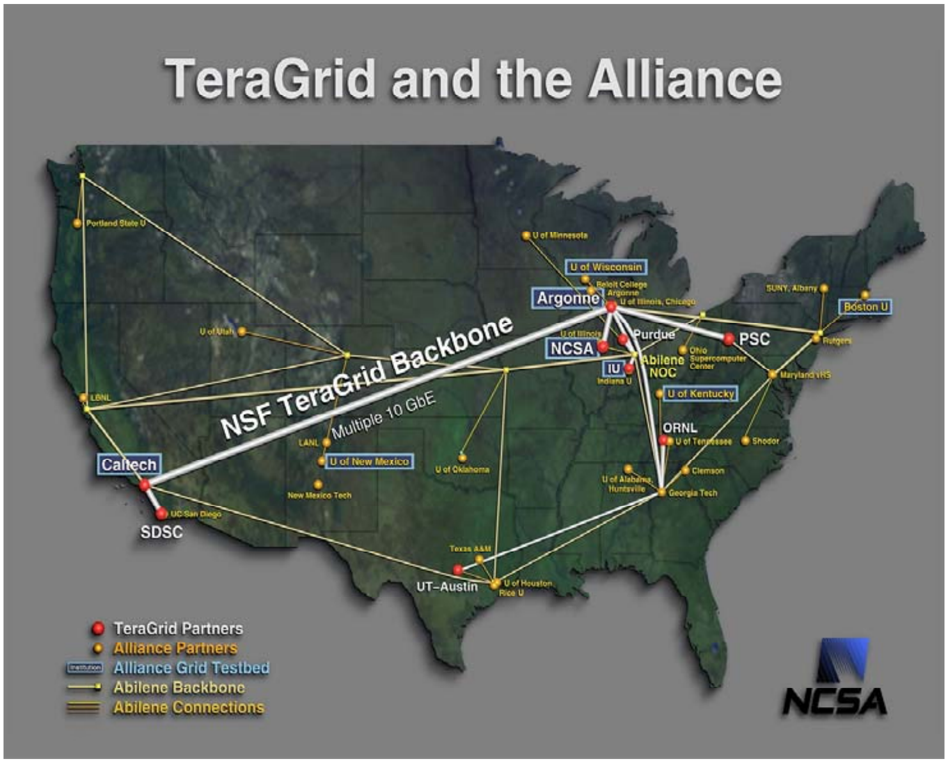

Figure 3: The NSF TeraGrid is designed to connect four primary super-computing sites in San Diego (SDSC), Pasadena (Caltech), Illinois (NCSA) and Chicago (Argonne) together via a 40 Gigabit Ethernet backbone. Additional partners (the Alliance) connect to the TeraGrid via the Abilene network.

2.4.2 Earth System Grid

ESG, shown in Figure 4, is a $15 million, 5 year DOE project to construct a data grid to address needs of the climate community. High-resolution, long-duration simulations performed with advanced DOE / NCAR climate models will produce tens of petabytes of output. To be useful, this output must be made available to global change impacts researchers nationwide, at national laboratories, universities, research laboratories, and other institutions around the world. To this end, the Earth System Grid, a virtual collaborative environment that links distributed centers, users, models, and data, has been created to provide scientists with virtual proximity to the distributed data and resources that they require to perform their research (www.earthsystemgrid.org). Initial focus of this project has been to build the necessary infrastructure required to build a data grid including data cataloging, metadata infrastructure, discovery capabilities and tools to provide data integration, sub-setting and sub-sampling of data streams.

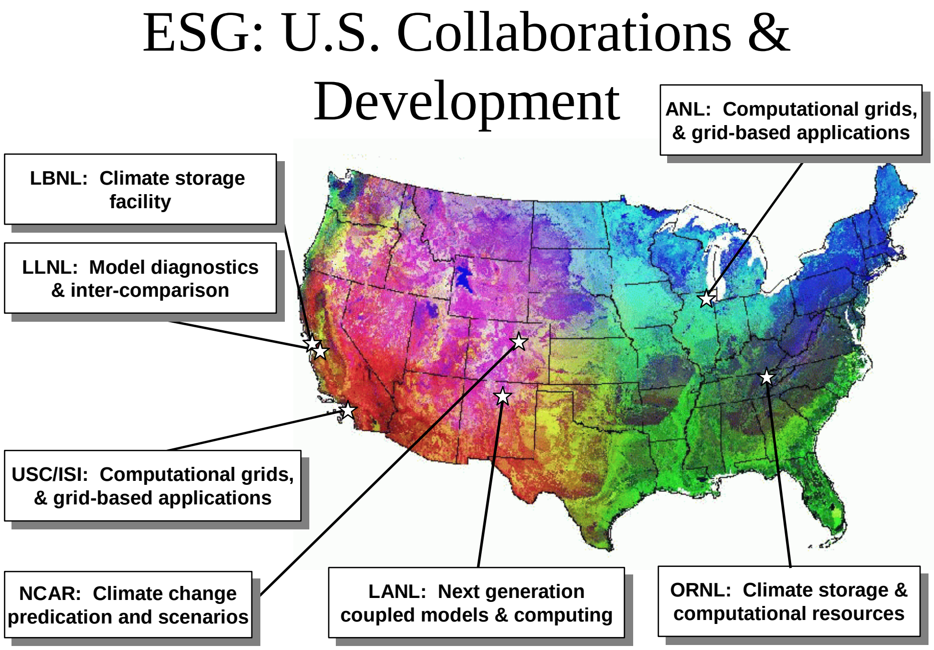

Figure 4: The Earth Science Grid is a data grid designed to support the sharing of results from climate simulations (courtesy of Don Middleton, NCAR).

2.4.3 LEAD

LEAD, illustrated in Figure 5 is an $11 million, five-year multi-university NSF funded project to "create an integrated, scalable cyber-infrastructure for meso-scale meteorology research and education." This project, led by the University of Oklahoma, proposes to develop an integrated, scalable framework for accessing, preparing, assimilating, predicting, managing, mining, analyzing, and displaying a broad array of meteorological and related information, independent of format and physical location. The focus of LEAD is to improve the capability to provide timely warnings of severe weather events by developing a dynamic computing and networking infrastructure required for on-demand detection, simulation and prediction of high-impact local weather such as thunderstorms (www.lead.ou.edu). In a LEAD scenario, for example, prior to a model run, all data relevant to the forecast would be discovered and gathered into a high-resolution data stream that would be integrated into the model. LEAD is being designed to potentially control select local observing systems, such as local radars, which could be directed by and take special observations for grid-enabled applications.

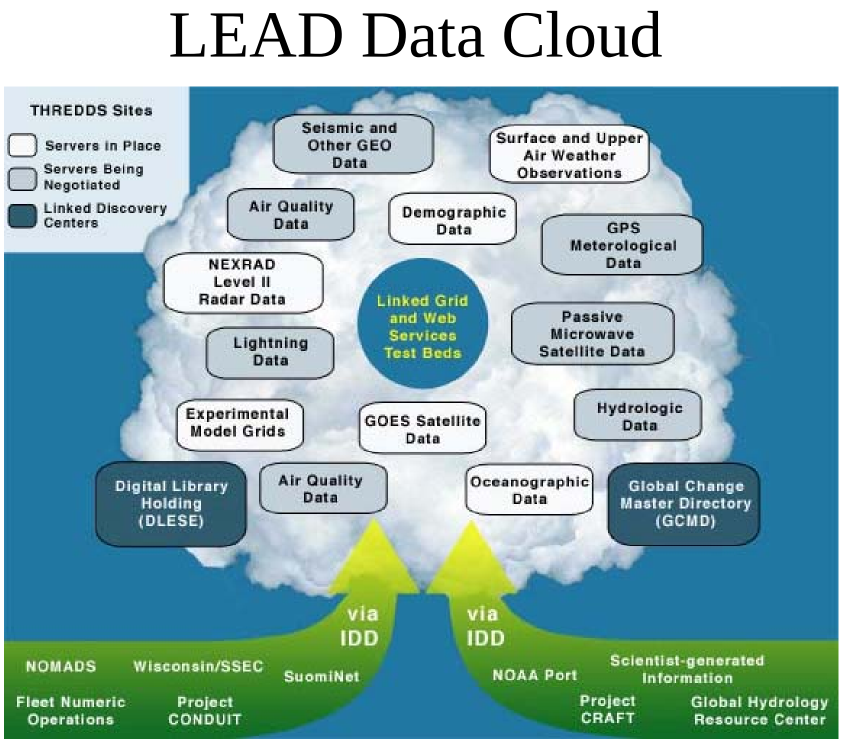

Figure 5: Data from disparate sources are integrated into the LEAD Cloud and then provided via grid and web mechanisms to applications and users that need them (courtesy of K. Droegemeier, University of Oklahoma).

3. Deploying Grids at NOAA

The NOAA Administrator, Vice Admiral Conrad Lautenbacher, describes the end goal of NOAA's data systems as a "fully wired, networked and integrated system that provides for data processing, distribution and archiving." The system for the creation, dissemination, distribution and use of these data are complex and involves many facilities in NOAA. Figure 6 illustrates many of the facilities that handle NOAA's data. The flow of data begins with the ingest of raw data from a host of instruments, remote sensors, and observing platforms including aircraft, satellites, radars, profilers, ships, and buoys. These raw data are used in weather forecast models, climate models, and for research applications. They are archived by NOAA's data centers, and saved for shorter periods at the local facilities where research or development is performed. To support this data flow, storage is required to hold the data, compute resources are required to create the data, and a network is required to move data to applications and users.

NOAA's IT infrastructure is currently a set of independent systems that have been built up over time to support its programs and requirements. Grid offers a capability to build an integrated, coordinated, cost effective IT resource to meet the needs of the entire organization. This section will describe how compute, data and service grids can be effectively built to meet the many scientific and IT challenges facing NOAA in the 21st century. The successful implementation of these grids relies on a robust network infrastructure, so discussion will begin with an examination of the NOAA network.

NOAA Laboratories, Centers and Forecast Offices

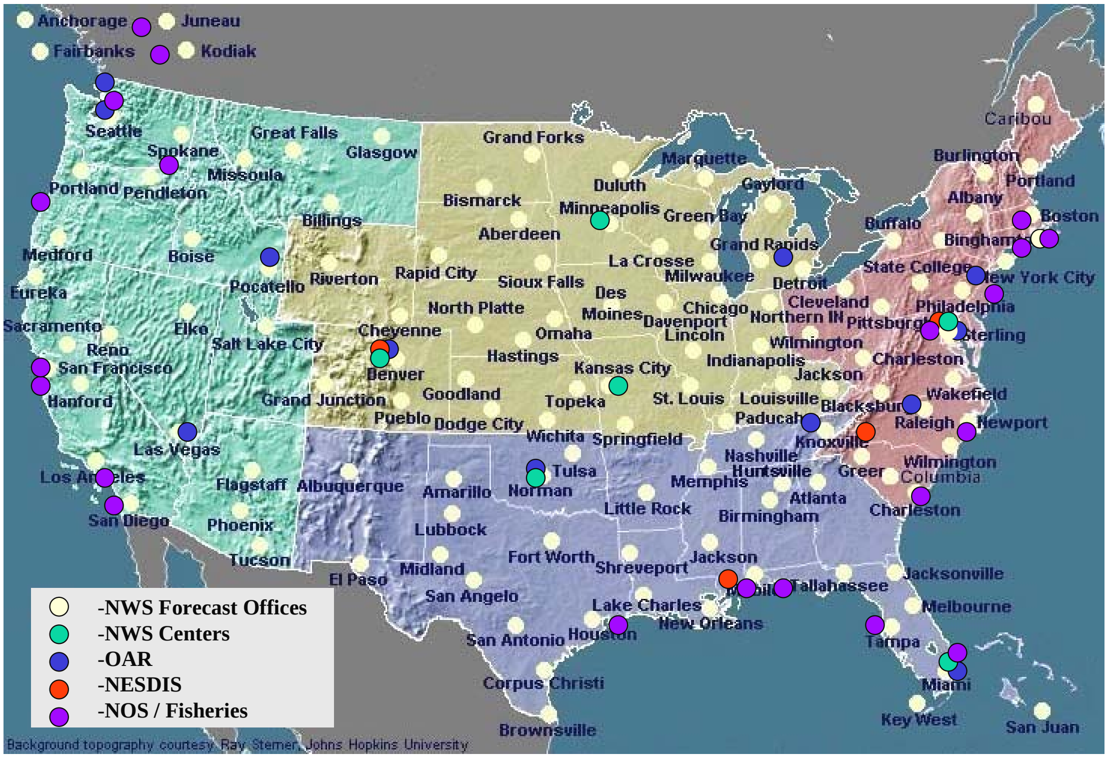

Figure 6: From local weather forecast offices to data centers and super-computing centers, NOAA facilities exist in every state in the nation. These sites all handle, manage, archive, or disseminate data within the organization and to agencies, universities, businesses and the general public.

3.1 The NOAA Network

The network is at the heart of NOAA's IT infrastructure; its existence is vital to the organization in being able to conduct research and to disseminate forecasts, warnings, advisories, and other information to federal agencies, businesses and the general public. The NOAA network extends to facilities located in every state in the union, and provides connections to collaborators at NASA, Navy, Air Force, FAA, universities, research laboratories, businesses and the general public. Data handled by NOAA are currently disseminated to points around the world. With expected collaborations in building the NOAA-backed Integrated Earth Observing System, the volume of data flowing between NOAA and many nations will significantly increase demands on the network.

3.1.1 The Current Network

The NOAA network includes the Satellite Broadcast Network (SBN) or NOAAPORT for broadband data distribution, the Internet, and many dedicated point-to-point links. The network, as illustrated in Figure 7, is functionally very effective, however, when considered as a single corporate solution it is a poorly organized, expensive, and inefficient resource. This is due to the growth of the network based on local and regional requirements driven by the distributed nature and geography of NOAA line offices, where inter-office coordination was not a high priority. Administration and management of such a diverse aggregation of network resources creates tremendous challenges in the areas of interoperability, reliability and security, and requires significant staff time at many locations to maintain the network. Janssen [2003] highlights many of these problems and proposes an integrated, coordinated, cost-effective, efficient NOAA network to replace or upgrade the existing network.

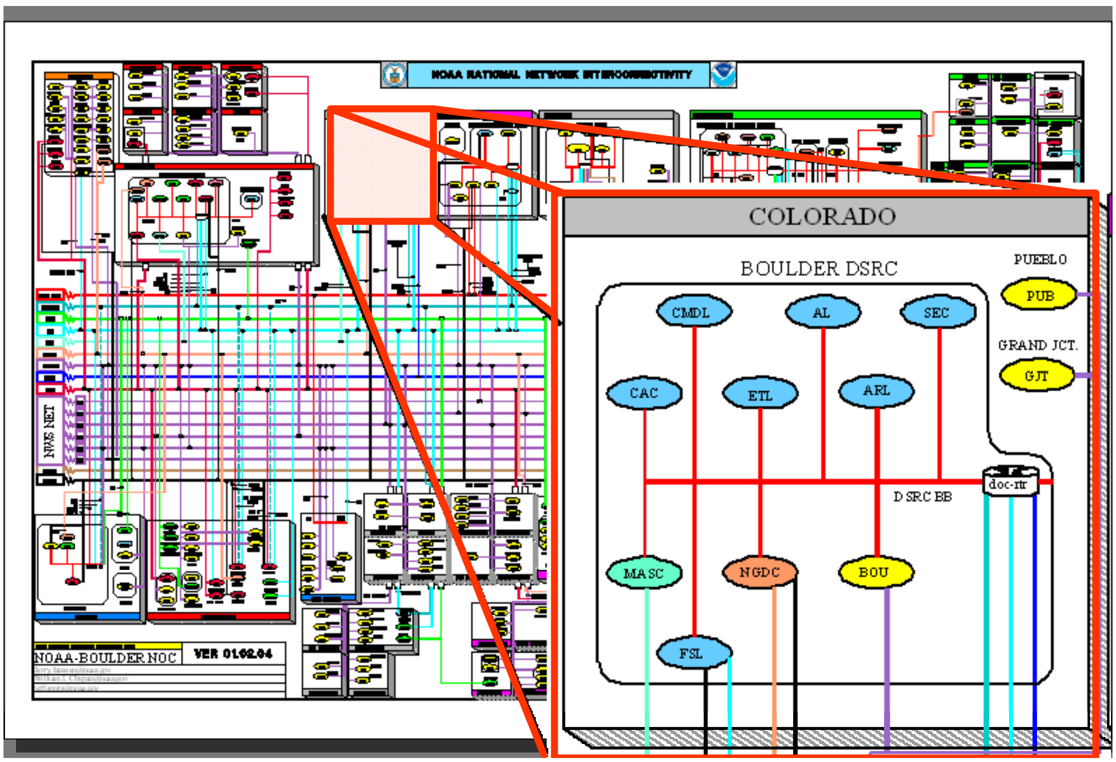

Figure 7: A diagram, developed by Janssen, describing current network connectivity between NOAA facilities. The zoomed portion illustrates network connections by NOAA laboratories at the Boulder site. Recognizing the problems associated with management, security, access, and administration of this network, upper management in NOAA consider this a significant problem that must be addressed.

NOAAPORT has been the main method of data distribution to NWS Weather Forecast Offices (WFOs) in support of AWIPS. However, bandwidth constraints have always limited NWS WFO's access to high-resolution data products. Recent efforts to free up bandwidth on NOAAPORT will help, but are seen as a short-term solution to the long-term problem of rapidly increasing data volumes. Increasingly the NWS and operations are using the Internet to access and distribute data. For example, the CRAFT program uses the Internet to distribute WSD-88D radar data to NWS WFO's, to NCDC for archiving, to Unidata for distribution to universities, and to other data users as required [Sandman, 2003]. The CRAFT program, now called the Integrated Radar Data Services (IraDS) (www.radarservices.org), has demonstrated that reliable, secure data transport can be accomplished for the NWS via the Internet, thus reducing the need for high-cost dedicated point-to-point links and paving the way toward an integrated, high speed network at NOAA.

3.1.2 National Lambda Rail

At the forefront of high-performance network technologies, is the emergence of national-scale fiber optic networks designed for research environments. The National Lambda Rail (NLR) network (www.nationallambdarail.org) is a consortium that is building such a fiber network to meet the research and development requirements for NOAA. NLR will utilize the ubiquitous Ethernet protocol over Dense Wave Division Multiplexing (DWDM) to create a scalable national network that supports both research and production environments simultaneously. DWDM uses multiple wavelengths of light (lambdas) to segregate unique networks as desired. Initial bandwidth will be available in one and ten-Gigabit Ethernet (GigE) segments, but is currently scalable to 40 x 10 GigE networks. The first segment of the NLR was "lit" in November 2003; the last segment of the first phase is scheduled for completion in December 2004 and will provide coast-to-coast lambda connectivity.

NOAA is fortunate to already have excellent connectivity to the primary partners and community network partnerships called GigaPoPs that are driving early development of NLR. NOAA sites in Boulder, Seattle, Miami and Washington D.C. already have NLR connectivity through GigaPops. Additional NOAA sites, including Princeton (GFDL), are exploring this opportunity. By utilizing dark fiber, network cost savings over typical wide area network services are substantial. For example, the cost to procure a 10 GigE link from Boulder to Denver, Colorado (~20 miles) from a network service provider is ~$330K per year, whereas the cost for the NOAA agencies in Boulder to join the NLR is $85K per year for a national (~10,000 mile) 10 GigE network. The attractive pricing, coupled with increased acceptance of the Internet by operational and research group, make the adoption of NLR by NOAA likely in the near future. This represents an important step toward building an efficient scalable distributed IT grid infrastructure at NOAA.

3.2 Compute Grids

Compute facilities are expensive to build and maintain, so being able to share cycles across the organization represents a more effective way to utilize IT resources. Compute grids within NOAA can be deployed at a local level or encompass computing facilities around the nation. Local compute grids can be deployed at a single site to make use of spare compute cycles that would not otherwise be utilized. These intra-site grids would complement the large supercomputing resources and would be useful for modest size computing tasks. One such grid at NOAA, developed by the NESDIS Office of Satellite Operations (OSO), is a small PC-based grid to test and verify the development of new data products [Guch 2004]. FSL is also constructing a small intra-division grid composed of desktop systems to conduct simulations and run applications during quiet periods.

Compute grids linking super-computing facilities at FSL, GFDL, and NCEP can also be used to smooth out system workloads between sites. Figure 8 illustrates the use of a grid-enabled meta-scheduler to interact with the queuing system at each HPCS site, to coordinate workflow between sites, to achieve better utilization of system resources, and to provide a high level of redundancy when required. For example, if compute facilities become unavailable at one site, the task can be migrated to another system on the grid. Grid schedulers can also be used to manage IT resources and to support resource reservation capabilities demanded by operational processes. Large-scale compute grids can also be used to solve grand challenge problems requiring access to significant IT resources that may not be available at a single site. Some applications that can benefit by the availability of such compute grids are further described.

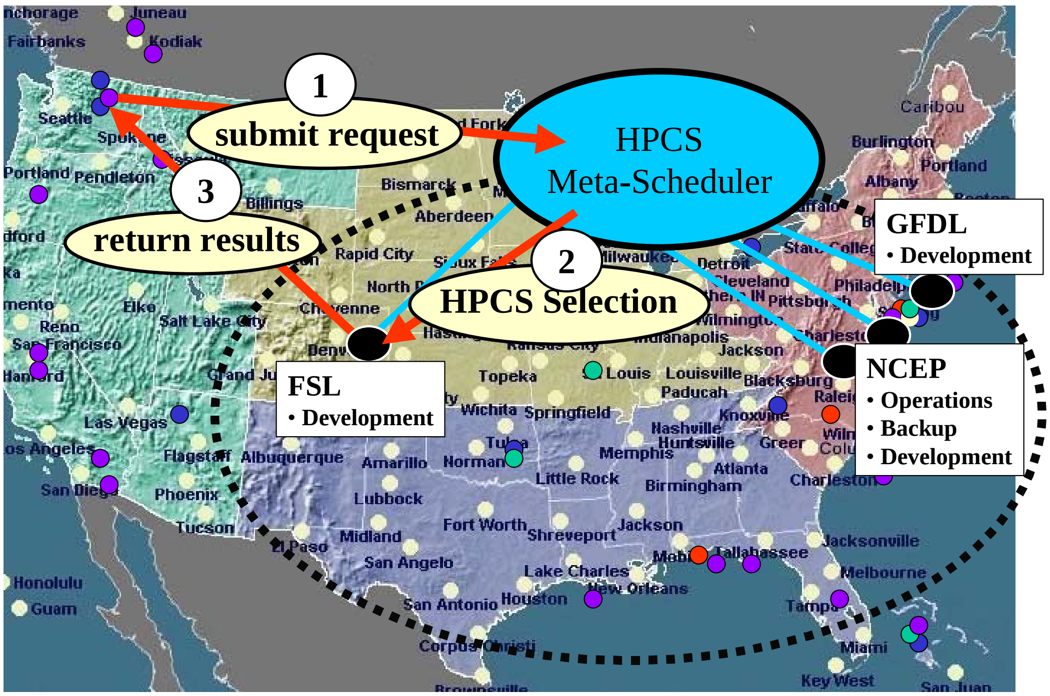

Figure 8: A NOAA compute grid to link super-computing sites within the organization. The workflow is managed by a meta scheduler as illustrated: (1) jobs are submitted to the HPCS meta-scheduler. (2), the scheduler interacts with each system's job queue to determine the best place to run the task, and (3) output results are returned to the host system.

3.2.1 Scientific Applications

NOAA's requirements for computing and data storage are growing rapidly in response to the increased flow and availability of data through its facilities to support (1) increasingly complex models, (2) new areas of research such as ensembles, data assimilation, and coupled models, and (3) the creation and integration of new high resolution datasets coming from next-generation observational platforms. More complex models stem from higher model resolutions, more accurate physical parameterizations, and the availability of higher density datasets. In the next decade for example, it is likely that regional climate models will be run on a "meso" scale. Since a doubling in model resolution requires a factor of 8 increase in the number of computations, moves to higher resolutions will be contingent on having the necessary resources available to run these applications.

Improvements in forecast accuracy are tied to better model initialization, data assimilation and model ensembles. Four-dimensional data assimilation techniques used to initialize models, offer promising research results but are hindered by the lack of the large compute resources required to run them. Ensembles combine the results of several to tens of instances of a model (or models) and/or perturbed initial conditions that are combined to achieve more accurate forecasts. An ensemble prediction system based on the Weather Research and Forecasting (WRF) model will be the next operational model in 2004. Since ensemble member runs are independent of one another, it would be feasible to have tens or even hundreds of ensemble components running across a NOAA computational grid.

3.2.2 *Verification*

As models become more complex, the number of data sources increase, and data density from observing platforms grow, verification will play an increasingly important role in both model development and observing systems evaluation. Improvements in forecast accuracy and increasing the speed new modeling and data assimilation techniques are transferred into operations is central to the NWS mission. The increased role of verification in model development is being demonstrated in the test and evaluation of WRF. A WRF ensemble is slated to become the next national-scale operational weather forecast model for NWS in October 2004. Prior to acceptance, WRF is undergoing an extensive series of tests identified in an NWS document called the WRF Test Plan [Siemens 2003]. Under this plan, a total of 1920 runs (idealized, platform, and retrospective) were run at Air Force Weather Agency (AFWA) and FSL to compare two variants to the current operational model. Additional testing and verification at NCEP is planned for the WRF ensemble prior to operational acceptance.

A second benefit of verification is its value in determining cost/benefit analyses of new or existing observing systems that cost NOAA hundreds of millions of dollars annually to build, deploy and maintain. For instance, data denial experiments can be run to quantify the value of existing observing systems. Simple scenarios, such as removing select RAOB observations, have been run to determine the effect on forecast accuracy [NAOS 2000]. More complex scenarios can be run to replace older expensive observing systems with a combination of multiple lower-cost alternatives. Verification can also be used to determine the value of deploying new or proposed observing systems. For example, the National Profiler Network (NPN) and National Radar Network (NRN) are slated for deployment in the next 5 years. Other systems such as the GPS-Met water vapor product may provide significant value if fully developed and deployed. Verification systems, such as the Real-Time Verification System (RTVS) [Mahoney, et al], can be used to compare model results and to provide verification statistics that help management and budget analysts determine which observing systems to develop and deploy.

To obtain accurate and valid verification statistics, extensive testing must be conducted over long time periods and will require significant IT resources. For example, many variants have been proposed as candidates for the operational WRF that could be evaluated resulting in tens of thousands of model runs, petabytes of data storage, and massive compute resources. Modest testing scenarios have been proposed for WRF in FY06 that will require 160% to 680% of FSL's

HPCS [Hart 2003], a system that is currently the 17th fastest supercomputer in the world (www.top500.org). This type of systematic test and verification can and should be applied to other models important to NOAA; however, a grid infrastructure must be built to coordinate and harness the resources required to run the tests.

3.2.3 Compute Grids to Support VOs

Figure 9 illustrates some examples of compute grids that could be created to support VOs within NOAA. An NWS grid could be designed to provide compute capabilities using unused cycles in the WFOs in support of local modeling efforts. WFO collaborations with NCEP could be used to deliver additional high-resolution data streams, and provide modeling support. A NESDIS product development VO could enable a richer collaboration between developers and users of new generation satellite data products. A Developmental Testbed Center (DTC) grid could be built to bring together research and operational aspects of WRF model testing, sharing both test results and new modeling developments. Finally, a HPCS grid could provide better integration and utilization of a supercomputing test and development facility to tackle grand challenge problems important to NOAA.

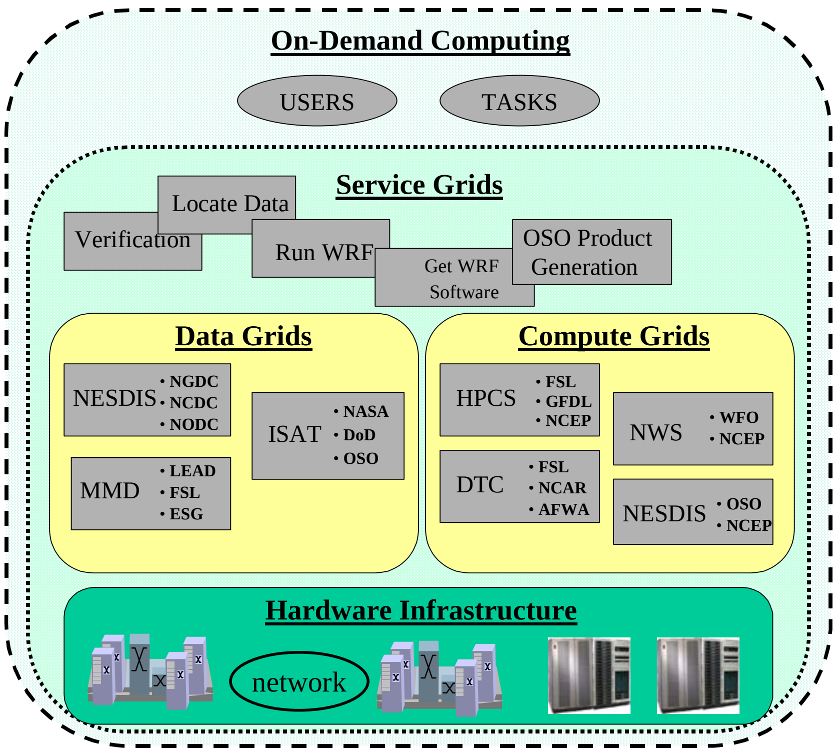

Figure 9: An illustration of how grids might be deployed to serve the needs of Virtual Organizations in NOAA and their collaborators. Critical services can be developed to permit better collaboration and coordination between work groups and across scientific disciplines. Several examples of Compute, data and service grids are shown that could be created to support VOs within NOAA.

3.3 Data Grids

Section 1 highlighted the expected enormous influx of data that must be managed by NOAA in the next decade. From building an infrastructure to support the Integrated Earth Observation

System, to handling expected growth in an increasingly dense set of observing systems, how to handle these data is likely the most compelling problem facing NOAA in the next decade. The volume of data expected in new high resolution observing platforms will stress NOAA's ability to archive, discover and access data in a meaningful way. As NOAA moves into a new era of even higher volume data, its IT infrastructure must be scalable, flexible and capable of meeting the rapidly changing requirements of the organization. Fundamentally, if data are not utilized, NOAA's sizable investment in the systems to provide those data streams is diminished.

The largest and most expensive data stream handled by NOAA is satellite data from GOES, POES and other remote-sensing platforms. These satellite data, operated by the NESDIS satellite operations facility currently distribute over 4TB / day for NOAA, NASA and DoD; its successful operation is vital to NOAA and the nation. In addition, numerous observational data are taken by NOAA, from RAOBS to precipitation gages, stream flow data, and snow cover. These data are in different formats, and different QC procedures often apply. Further, new national observing systems including profiler, GPS water, and radar are being proposed or slated to become available. These data represent numerous "stove pipe" developments that make it difficult to coordinate their use across scientific disciplines.

The NOAA Observing System Architecture (NOSA), illustrated in Figure 10, is building a system to define, coordinate and integrate NOAA data sources; currently NOAA has found that it has 102 separate observing systems measuring 521 different environmental parameters. NOSA is billed as a program to both identify observing systems and provide a single integrated system to deliver information "on demand" to defense, commercial, civilian and private sectors at the national and international levels. So far, NOSA has focused on identifying NOAA's observing systems as a means to highlight gaps and duplication in current observing systems. However, delivery systems will need to be developed to locate, access and integrate data into systems and users that require them. These systems will need to be distributed in nature, in order to account for data located in repositories around the nation and the world. Grid mechanisms for data access, discovery and integration can provide such a distributed delivery system, improve data utilization and address the future needs of NOAA.

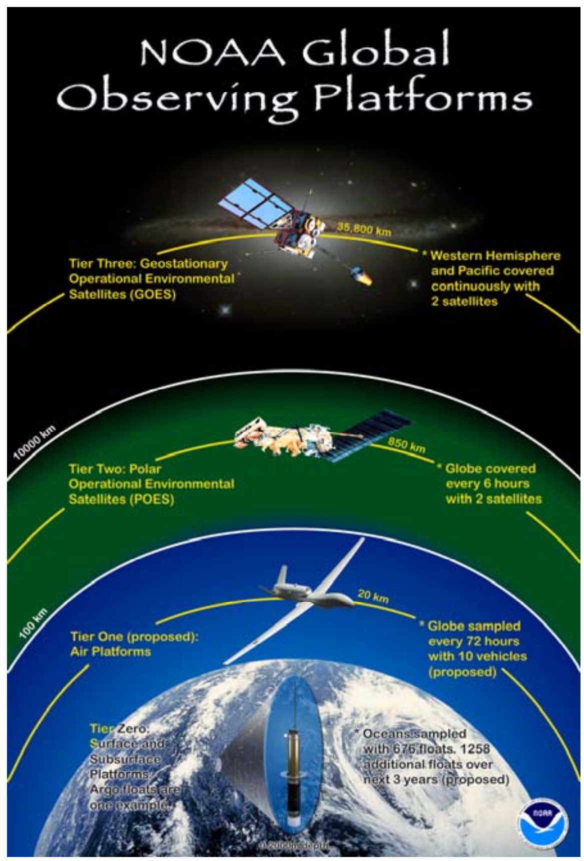

Figure 10: NOAA is conducting an inventory of all of its observing systems and creating the NOAA Observing System Architecture to document and integrate them into a single framework (Courtesy of *NOAA Magazine*, April 2004).

3.3.1 *Leveraging Existing NOAA Programs*

Data grids can provide an "umbrella" technology to link existing NOAA data infrastructure into a more coherent integrated system. Programs such as the NOAA Operational Archive and Distribution System (NOMADS) project and the Comprehensive Large Array-data Stewardship System (CLASS) are good examples of existing systems that can be used as a foundation for integrated data grids. NOMADS is a distributed data archive system that provides mechanisms to access and integrate model output and other data stored in distributed repositories. NOMADS enables the sharing and inter-comparing of model results and is a major collaborative effort, spanning multiple government agencies and academic institutions. The data available under the NOMADS framework include model input and Numerical Weather Prediction (NWP) gridded output from NCEP, and Global Climate Models (GCM) and simulations from GFDL and other leading institutions around the world (www.ncdc.noaa.gov/oa/climate/nomads/nomads.html). CLASS is an online operational system that has been designed and funded to archive and distribute NOAA's observational data; by 2015 CLASS is expected to be archiving over 1,200 TeraBytes of data per year. Dual redundant CLASS systems provide access to observational data via the web, and the project plans to offer enhanced database management, search, order, browse and subsetting techniques to users and applications (www.class.ncdc.noaa.gov).

While the NOMADS and CLASS systems represent effective means to obtain observational data within NOAA, more robust mechanisms for the discovery, access, and integration of data will be required to effectively utilize all data across the organization and across scientific disciplines. In a data grid scheme, NOMADS and CLASS can be viewed as one of many potential data streams that also include local data sets (e.g. model runs, experimental datasets), proprietary data (e.g. airline data) and other data important to NOAA. These data can be made available via the grid using web-services, discovery languages such as XML, and grid interfaces deployed by a growing number of commercial DBMS systems from Oracle and others.

The basis from which data grids derive significant benefits over existing systems stem from grids underlying use of web technologies to permit dynamic discovery, access, and integration of distributed data. For example, the construction of web-accessible metadata catalogues can be utilized to identify data stored in the web-friendly XML format. Web and grid services can be used to dynamically locate meta-data catalogues and deliver data to the end-user or application upon request. Current static methods used to obtain data via CONDUIT, OpenDAP, and others require apriori knowledge of data sources and agreements between the data provider and user (client / server).

To permit dynamic access and discovery of data, the development of generalized data catalogues that identify data sources, history, and information about the data are very important. Tremendous work within NOAA has already been done to catalogue and create metadata mechanisms. This is a significant and essential effort necessary to providing generalized grid access to data. The ESG project, for example, has spent significant effort to define sufficient information to categorize climate data, while building grid mechanisms to access, integrate and subset the data. The LEAD project proposes to extend this work by exploring dynamic access and discovery of data in quasi-operational environments.

3.3.2 New Data Storage Strategies

A fundamental problem with the current data is an inherent inefficiency in how data are stored, accessed and used. Data delivery mechanisms such as CONDUIT and OpenDAP provide valuable means to distribute data and are useful in operational settings where data utilization may be high. However, in some research and operational settings where data utilization may be low, storing local copies of data available elsewhere is not cost effective.

Slow networks exacerbate the need to store local copies of data that are often never accessed. For example, FSL stores copies of operational Eta model runs that are also being stored at NCEP and NCDC and likely many other sites in NOAA. FSL saves over 200 Gbytes of its real-time data to its Mass Storage System daily and only a fraction of it is ever requested. This is a typical pattern of usage mirrored across sites in NOAA. The cost of archiving this data mounts as tapes must be purchased, offsite storage must be located and maintained, older data must be migrated to newer storage media, and operator time is required to handle the data. With a faster network, and better methods to locate data, the requirement that it be stored locally is reduced. Ideally, data should be stored in a few places for redundancy, and obtained from these sites when required (on-demand). Finally, it may also be useful to not archive some data, but store the recipe used to create it instead. Then when the data is required, the relevant programs can be run to produce it.

3.3.3 Potential Data Grids at NOAA

Figure 9 illustrates several data grids that could be built for VOs in NOAA to satisfy operational and research requirements. For example, a NESDIS grid could be built to integrate deep storage archival systems under a single service-oriented infrastructure. Users requiring access to archival data could locate and access it when required using grid file transfer mechanisms. Another grid, the Integrated SATellite Data Grid (ISAT) could be similarly available to producers and consumers of this data, allowing new product developments at NESDIS, to be more easily tested in next generation models at NCEP. A prototype grid is already being used at NESDIS to test new products, but systems are currently limited to desktop PCs as stated in Section 3.2. Access to large storage facilities could provide a more efficient means to test new products required by new generation high-resolution data streams. A Meso-scale Modeling Data (MMD) grid could be designed for model development described in Section 3.2.2. Collaborators could share test results and access information about what tests have been run thus reducing duplication.

3.4 Service Grids

NOAA continues to evolve into an organization that is increasingly run like a business: a responsive, flexible, service-oriented organization that can adapt to changing requirements and demands. Services grids can map directly into functions such as delivery of data, generation of data products, and on-demand computing requirements for operational activities within NOAA. Figure 9 illustrates the tiered relationship between tasks available in a service grid infrastructure, the data and compute grids where they will run, and the hardware required to run them.

Service Grids represent the future of Grid, but will likely not be sufficiently mature for 2-5 years. These grids are expected to amplify the separation between QoS, on-time delivery of products or data and the hardware and software systems required to accomplish the tasks. The business community is already moving away from owning and operating their own IT infrastructures, choosing instead to contract these operations out. For example, the fastest growing and largest part of IBM is its Global Services sector. IBM Global Services provides both software and hardware infrastructure support for their customers. IBM offers service contracts to many customers including financial institutions, insurance companies and foreign governments. Contracts are often specified in terms of QoS (reliability, downtime) and Global Services designs an end-to-end system to meet those requirements.

In NOAA, the NCEP HPC contract is structured in the same way. The NCEP HPCS system is owned by IBM and housed at IBM controlled facilities. The NCEP contract states specific QoS metrics for performance, and delivery of products that are critical to the mission of the NWS. IBM must meet those guarantees within the bounds of the agreement, but is free to upgrade or dedicate more or less resources as they see fit. This decoupling of IT from mission critical requirements represents a trend that is growing in the commercial sector, and is becoming more common within NOAA. In the future, both hardware and software contracts could be let for compute, data, archival, and network services as a cost savings measure. This would allow NOAA to focus on core activities and less on the IT infrastructure required to accomplish its goals.

4. Building a Grid Infrastructure at NOAA

4.1 Challenges to Building Grids

There are several challenges that must be addressed in building Grids at NOAA. First, there is the need to build an efficient network infrastructure to permit distributed computing and integrated data resource capabilities to be effectively utilized via the Internet. The backbone for an integrated network will likely be NLR. For NOAA, the NLR creates an opportunity to link geographically diverse offices and IT resources via major regional network centers (GigaPoPs) as if they were part of one local area network. While NLR is a good first step, NOAA needs to move away from simply upgrading individual links and design an integrated resource that permits efficient data flow to all NOAA sites. For example, network links to WFOs should be upgraded so forecasters can receive the high-resolution model data they need to produce more accurate forecasts. The cost in building such a network will be mitigated by the design of a more efficient and integrated network, and the benefit to the organization of sharing resources and data more effectively. Additionally, prioritization of upgrades will be required.

Second, there is the need to build a secure grid infrastructure that will minimize vulnerabilities, and maximize reliability, robustness, and flexibility. The NOAA network security model requires that firewalls be used to protect IT assets from electronic attack by permitting very restricted access by “external” systems (outside the firewall) to “internal” systems (behind the firewall). Virtual Private Networks (VPNs) are created dynamically to establish a “secure” tunnel between external and internal systems in which all network traffic is encrypted. To overcome increasingly sophisticated cyber attacks, NOAA will likely begin using token-based authentication (one time passwords) in the near future.

This NOAA network security model does not need to be changed to support grids. The use of network “secure zones” can be created by dynamically creating VPNs to permit grid traffic to pass through firewalls. These secure zones will permit “internal NOAA” traffic from trusted sites to share grid resources. Some performance issues will need to be explored including the impact of VPN encryption and the ability to handle large data flows through NOAA firewalls. Grid tools provide the methods to achieve highly secure grids, but accounts, permissions, and certificate trust relationships must also be created and maintained. NOAA will need to continue building a Certificate Authority (CA) to handle trusted third party authentication between hosts. [James 2003] In addition, a central coordination group will need to be formed for each VO to determine an acceptable model of trust between members.

Third, there is a need to build software systems that emphasize portability and better access to IT resources. As NOAA stresses the increasing importance of sharing IT resources across the organization, the ability to port codes between systems for development and use will become a critical requirement. A recent NOAA HPC study detailed the benefits of combining HPCS contracts and will lead to a single procurement in FY05 to support multiple constituencies. As a result, significant efforts will be made to port codes across HPCS systems in the next few years, a process that will uncover bugs, improve reliability of the applications, and improve systems utilization. Additionally, the increased reliance on the web as a medium for data access and discovery will spawn a new generation of grid enabled applications called portals. A grid portal is a web-based application designed to connect the user with grid resources necessary to accomplish desired tasks. One might have a data portal to access climate data, a compute portal to run the WRF model, and a monitoring portal to monitor an operational run. As IT resources become easier to access and codes become more portable, better utilization of software and data is assured and new interdisciplinary studies become possible. Part of the challenge will be to design applications that are sufficiently flexible to be used in ways now unimagined.

Finally, there is a need to engage management in the importance of building an IT infrastructure that emphasizes common requirements of the organization, rather than the needs of individual constituencies. The driving factor to improve collaboration and to share IT resources within NOAA is cost reduction. The recently enacted Planning, Programming, Budgeting and Execution System (PPBES) defines a set of crosscutting programs that span the organization for the purpose of reducing duplication, focusing efforts, and providing better oversight. The creation of VOs, enabled by a grid infrastructure, will permit these requirements to map directly into IT resources and enhance collaborations necessary to meet the given tasks.

4.2 Current Efforts

FSL is currently engaged in two projects to explore the Grid and how it can be utilized at NOAA. First, FSL is collaborating with OAR's Pacific Marine Environment Laboratory (PMEL) and the Geophysical Fluid Dynamics Laboratory (GFDL) in building a simple prototype NOAA grid. This work, funded by NOAA High Performance Computing and Communications (HPCC), explores the ability to construct a working grid, handle administrative and technical issues including resource allocation, security, and job scheduling, and run real scientific applications across wide area networks. This project, titled "Building a Prototype NOAA Grid", plans to initially use several coupled modeling applications at GFDL and PMEL and builds on a previously funded HPCC proposal titled "Running coupled models over the grid."

Second, FSL is developing a prototype grid portal to be used in support of the development, test and verification of the Weather Research and Forecast (WRF) model. Model test and verification is a tremendously complicated procedure that requires porting codes to multiple platforms, testing for correctness, building scripts to insure tests are run correctly, and examining the results. Development of a grid portal is being explored as a way to simplify testing, share research results and compute and data resources with our collaborators in NOAA, the National Center for Atmospheric Research (NCAR), and the DoD.

FSL is also an unfunded collaborator on the LEAD project; LEAD hopes to use FSL's quasi-operational capabilities as a test bed for near real time forecasting. FSL hopes to leverage LEAD expertise in building and testing applications using grid-based tools for data storage, access and discovery. FSL is also collaborating with the ESG project to share mechanisms for data discovery and access, formats and standards where such commonality exists. Finally, discussion have occurred with FNMOC about mutual interests in building organizational grids and in using the WRF model for grid tests.

4.3 Next Steps

As NOAA begins to take advantage of the enormous potential for grids, it will want to be positioned behind the bleeding edge of innovation, able to take advantage of new developments as they show promise but only use those technologies that are proven, stable and secure. To build a grid infrastructure, NOAA needs to seek partnerships with all its line offices and leverage existing programs such as CLASS and NOMADS. FSL, with its mission to explore and develop new capabilities required by NOAA, is well positioned to lead the research and development of grid at NOAA. FSL operates a shared compute facility used by many research laboratories and line offices within NOAA. There is also very good network connectivity to the Boulder Laboratories, that will permit others to access resources and to test grid applications using the Abilene network today, and the NLR network in the near future. In addition, FSL and NESDIS have a long history in building data systems for operational and research constituents that can be tapped to build distributed data management systems based on grid. As these technologies are developed, they can be transferred to operational communities at NWS, NESDIS and other NOAA line offices.

Development of grids must proceed in three areas. The single most important task in building such a grid infrastructure is to build a network that is reliable, secure, flexible, and sufficiently fast to permit sharing of IT resources. A comprehensive analysis of data flow through NOAA facilities will help determine needs and priorities in upgrading the network. High-level network discussions are ongoing but considerations should also incorporate supporting grids and increased demand-based data access into the design. A second area of development is the evolutionary construction of data and compute grids that integrates IT resources, provides new generations of web-based applications and offers generalized data availability to NOAA and its constituents. Initial efforts will continue at FSL to demonstrate the feasibility and benefits of grid-enabled applications and portals. As grids become more mature, grid computing environments and portal developments will be increasingly used to run interoperable models, locate data and analyze results. The third area of development is the construction of service grids, which will build on the success of data and compute grids deployed at NOAA. Generalized services can be developed to first support data access, discovery and retrieval. Additional services can be developed to support operational processes at NWS and NESDIS, and to simplify the research and development of next generation models that utilize high-resolution data streams expected in the next decade at NOAA.

5. Conclusion

This paper highlights the need to build a grid infrastructure to meet the challenges facing NOAA in the 21st century. Given the enormous expected demands for data, and increased size and density of observational systems, current systems will not be scalable for future needs without incurring enormous costs. NOAA needs integrated data systems capable of handling a huge increase in data volumes from expected launches of GOES-R, NPOESS, new observing systems being proposed or developed, and to meet requirements of the Integrated Earth Observation System. Further, NOAA must continue moving toward the use of shared compute resources in order to reduce costs and improve systems utilization, to support new scientific challenges and to run and verify increasingly complex models using next generation high-density data streams. Finally, NOAA needs a fast, well-managed network capable of meeting the needs of the organization: to efficiently distribute data to users, to provide secure access to IT resources, and be sufficiently adaptable and scalable to meet unanticipated needs in the future.

Just as the NOAA Observing System Architecture (NOSA) was developed to define, coordinate and integrate data sources, NOAA needs to build an IT counterpart (NITSA) to define, coordinate and integrate its IT resources. We propose the construction of a grid infrastructure, centered around a fast network, which will permit resources including computer, data storage, software systems and services to be managed and shared more effectively. All three types of grids can be utilized at NOAA: compute grids to provide access to under-utilized cycles and to link super-computing centers, data grids to promote sharing and better utilization of data, and service grids to provide the reliability and redundancy demanded by operational processes. Grid also provides robust mechanisms for data access, discovery and integration that can transform static methods of data distribution into dynamic demand driven delivery systems and thereby reduce the need to store redundant copies of data. We have also described a process by which secure grids can be built using token-based authentication, VPN communications between secure NOAA systems, and restricted external access through heavily monitored firewalls.

NOAA currently spends hundreds of millions of dollars on observing systems, supercomputing facilities, data centers, and the network. The development of a NITSA, enabled by Grid, can potentially reap huge cost savings for the organization. We described the value of building an efficient, scalable, secure and integrated network. It is not sufficient or cost effective to continue to upgrade select point-to-point network links. The entire network needs to be examined and an integrated, managed, coordinated network resource must be created that is an asset to the entire organization. We also described the value of building an grid infrastructure to coordinate access to the IT resources necessary to handle increasing volumes of data, to develop and run next generation models and to evaluate the cost/benefit of new and existing observing systems. Further, as NOAA stresses the increased importance of sharing IT resources in the interest of cost savings, code interoperability across HPC systems at GFDL, FSL and NCEP has become a critical requirement. It is not enough to simply achieve code portability, however; improved usability of these codes is also required. Grids provide a supportive distributed computing environment in which to run portable codes and simplifies user access to shared system resources. Finally, we described the importance of VOs (Virtual Organizations) to improve collaborations between organizational units, to improve efficiency and to reduce costs. Grids support the creation of VOs, which can map IT resources directly into cross cutting programs identified by NOAA's PPBES management plan.

Acknowledgements

The authors wish to thank Dan Schaffer and Jacques Middlecoff from FSL for their valuable feedback on ideas and research discussed in this paper. Paul Hamer and Chris MacDermaid from FSL were helpful in describing NOAA's data systems and offering ideas about how NOAA might handle high data volumes in the future. Jerry Janssen from OAR provided valuable information on the NOAA network and administrative plans to upgrade it. We also thank NOAAs High Performance Computing and Communications (HPCC) for sponsoring the annual NOAA TECH conferences. This forum has been a valuable way to learn about new technologies and developments at NOAA. The NESDIS Data Users and GOES-R conferences underscored the many challenges facing NOAA in the next decade. Finally, we thank the developers of the many excellent NOAA web pages that provided valuable information about the many exciting programs in the organization.

References

1. Atkinson, M., Chervenal, A., Kunszt, P., Narang, I., Paton, N., Pearson, D., Shoshani, A., and P. Watson, 2004, Data Access and Integration, *The Grid 2: A Blueprint for a New Computing Infrastructure*, pp391-430, Elsevier, ISBN: 1-55860-933-4.
2. Droegemeier, K.K., K. Kelleher, T. Crum, J.J. Levit, S.A. Del Greco, L. Miller, C. Sinclair, M. Benner, D.W. Fulker, and H. Edmon, 2002: Project CRAFT: A test bed for demonstrating the real time acquisition and archival of WSR-88D Level II data.

Preprints, 18th Int. Conf. on Interactive Information Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology., 13-17 January, Amer. Meteor. Soc., Orlando, Florida, 136-139.

3. Foster, I., C. Kesselman, J. Nick, S. Tuecke, The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration, Open Grid Service Infrastructure Working Group, Global Grid Forum, June 22, 2002 http://www.Globus.org/research/papers.html#OGSA.
4. Foster, I., Kesselman, C., Teucke, S., 2001, The Anatomy of the Grid: Enabling Scalable Virtual Organizations, International Journal of Super-Computing Applications. http://www.Globus.org/research/papers/anatomy.pdf .
5. Guch, I., 2003: Harnessing the Spare Computing Power of NOAA Office PCs for Improved Satellite Data Processing and Technology Transition, NOAA Tech Conference, October 2003, http://www.noaatech2004.noaa.gov/abstracts/ab42_guch.html.
6. Hart, L., Govett, M., Doney, M. and J. Middlecoff, 2003: A Discussion of the Role of the Boulder High Performance Computer, NOAA Internal Document to Carl Staton 12/29/2003.
7. James, R., 2003, NOAA Public Key Infrastructure (PKI) Pilot Project, NOAA Tech 2004, http://www.noaatech2004.noaa.gov/abstracts/ab16_james.html.
8. Janssen, J., 2003, The NOAA Network Baseline, Talk given at NOAA Tech 2004, http://www.noaatech2004.noaa.gov/abstracts/ab83_janssen.html.
9. Mahoney, J.L., Judy K. Henderson, Barbara G. Brown, Joan E. Hart, Andrew Loughe, Christopher Fischer, and Beth Sigren: 2002. The Real-Time Verification System (RTVS) and its application to aviation weather forecasts. Preprints. 10th Conference on Aviation, Range, and Aerospace Meteorology, Amer. Met. Soc., 13-16 May, Portland, OR.
10. NAOS Report, 2000, Radiosonde/MDCRS Evaluations and Analyses, Findings, Conclusions, and Next Actions, Alexander MacDonald – chair of NAOS Council.
11. NOAA's FY2003-2008 Strategic Plan, 2003, http://www.spo.noaa.gov/strplan.htm.
12. Sandman, T., and K. Kelleher, 2003: Real Time Dissemination of WSR-88 Radar Data over Internet 2, NOAA Tech Conference, November 2004, http://www.noaatech2004.noaa.gov/abstracts/ab45_sandman.html.
13. Siemens, N., 2003: The NCEP WRF System Test Plan.
14. Tabb, L., 2003, Grid Computing in Financial Markets: Moving Beyond Compute-Intensive Applications, The Tabb Group, LLC, Westborough, MA.
15. Wladawsky-Berger, I., 2004, The Industrial Imperative, *The Grid 2: A Blueprint for a New Computing Infrastructure*, pp25-34, Elsevier, ISBN: 1-55860-933-4.
16. Wolfe,D. and S. Gutman, 2000, Development of the NOAA/ERL Ground-Based Water Vapor Demonstration Network: Design and Initial Results, *J. Atmos. Ocean. Technol.*, **17, 426-440**.